White Paper Report

Report ID: 105636

Application Number: HJ-50085-12

Project Director: Michael Cuthbert (cuthbert@mit.edu)

Institution: Massachusetts Institute of Technology

Reporting Period: 2/1/2012-8/31/2014

Report Due: 11/30/2014

Date Submitted: 12/1/2014

ELVIS

Electronic Locator of Vertical Interval Successions:

The first large data-driven research project on musical style

Digging into Data Challenge, 2012

Final report

August 5, 2014

Julie E. Cumming, McGill

Frauke Jürgensen, Aberdeen

Michael Scott Cuthbert, MIT

TABLE OF CONTENTS

Part I: THE WHOLE TEAM

1. Introductory summary: Structure and achievements

The ELVIS project had three locally based teams (in Canada, at McGill, in Scotland, at Aberdeen, and in New England, USA, divided between MIT, the lead, and Yale), each of which focused on a different aspect of the overall research program: using computers to understand musical style. The central unifying concept of the ELVIS project was to study counterpoint: the way combinations of voices in polyphonic music (e.g. the soprano and bass voices in a hymn, or the viola and cello in a string quartet, as well as combinations of more than two voices) interact: i.e. what are the permissible vertical intervals (notes from two voices sounding at the same time) for a particular period, genre, or style. These vertical intervals, connected by melodic motions in individual voices, constitute Vertical Interval Successions. In more modern terms, this could be described as harmonic progressions of chords, but what made ELVIS particularly flexible was its ability to bridge the gap to earlier, contrapuntally-conceived music by using the diad (a two-note combination) rather than the triad (a combination of three notes in particular arrangements) as a basis (since triads and beyond may be expressed as sums of diads).

Our research required the acquisition and organization of large amounts of symbolically-encoded music now available, in order to be able to use computers to study it. Existing data, while numerous, were somewhat messy, with many duplications, errors, and gaps in certain areas of music history, so one task was the consolidation and cleaning-up of the data both by hand and with newly developed error-correction software.

Altogether, the ELVIS project has enabled not only the consolidation of data and toolsets, but the creation of concrete research output on a previously difficult level. The resulting databank and tools available through the main website at McGill University will prove an invaluable resource to musicologists in this field in the years to come. Together the three teams made significant advances on a variety of different fronts; the interaction among the teams was essential to the success of the research.

The McGill team focused on creating tools and repositories of data accessible to researchers in musicology and music theory without significant computer skills. Their major contributions are:

- The ELVIS database, a repository of over 6000 compositions in symbolic notation (searchable by computer), with metadata making it possible tailor a search with a great deal of precision. For example: you can search the database for a particular genre, from a particular period, from a particular country with a specified number of voices, e.g. madrigals composed in Italy between 1540 and 1570 with four voices.

- VIS: an open access web-based software with a graphical user interface that allows users to upload files containing music in symbolic notation (from the ELVIS database or elsewhere), search for Vertical Interval Successions, and represent findings in a variety of ways: in a table form, as a graph, or in an annotated score output. VIS is built on music21, which was created and developed by the MIT team.
- A website with links to the database and software, as well as additional documentation on the project, including an API for VIS.
- Research projects on style change, contrapuntal repetition, the connections between music theory and repertoire and on sonification of music data.

The Aberdeen team focused their research on several different projects using information about Vertical Interval Successions in specific repertoires. They compiled several useful new datasets of music in symbolic notation: further pieces by Palestrina, Monteverdi, and Gesualdo. Unlike the other parts of the team, they used the PerlHumdrum toolkit for Symbolic Music Processing. Humdrum (developed by David Huron) was the first major music-analysis software. Knopke and Jürgensen expanded it into the PerlHumdrum toolkit, which they continued to develop during the grant period. This allowed them to use some different kinds of search strategies; their focus was more on research results than on creation of tools for other scholars. Major achievements include:

- Techniques for identifying repeated chunks of counterpoint (n-grams or contrapuntal modules) where the match is not exact, using suffix arrays. They tested this against a corpus of contrapuntal analyses by Peter Schubert, a member of the McGill team (Schubert 2008).
- Study of historical tunings. Certain Vertical Interval Successions sound good or bad depending on the tuning system in place; this makes it possible to propose the probable tuning system for different repertories using different vertical intervals.
- The historical listener project. Dissonance and consonance are central issues in style classification; this project measured the amounts of dissonance in two contrasting repertoires and then asked listeners to rate the repertoires in terms of their dissonance.

The MIT/Yale team focused their research on the improvement and development of music21, a Python toolkit designed to query music in symbolic notation and on research on the chordal language of common practice music ca. 1700–1900 using music21. Major achievements include:

- Making music21 faster and better able to handle longer pieces of music and large corpora through multiprocessing and distributed computing.
- Developing tools that allow the user to select specific vertical intervals within a specific time frame for analysis.

- Developing tools for improving optical music recognition, in order to generate more searchable music in symbolic notation.
- Creating two major repositories of music: at Yale, a corpus of 5,000 large works from the common-practice period converted from MIDI files; at MIT, a curated collection of over 1,000 pieces of fourteenth-century music. Additionally, new parsers were written that expanded the range of pieces that could be read, including Rob C. Wegman's collection of hundreds of complete Renaissance Masses.
- Comparisons of the taught rules for tonal, chord behaviour and voicings with the statistical evidence from the repertory, often questioning or invalidating the received teachings.
- Research projects on fourteenth-century music, similarity and chordal progressions, alphabet reduction in chordal space, and the works of Bach and the Sacred Harp repertory.

2. People involved in the project

Canada (McGill): Julie E. Cumming, PI (for Canada, and for the whole project).
Co-investigators at McGill:

Ichiro Fujinaga, Music Technology
Cynthia Leive, Head Librarian of the Marvin Duchow Music Library
René Rusch, Peter Schubert, and Jonathan Wild, Music Theory.

McGill students who worked on ELVIS
Musicology

Remi Chiu (PhD; now in tenure-track position in the US)
Catherine Motuz (PhD): Database, blog
Jacob Sagrans (PhD): Database, corpora
Jane Hatter (PhD): Coordination, scheduling
Daniel Donnelly (PhD): Database
Rory O'Connor (MA Musicology): Database

Music Theory

Christopher Antila (MA) Lead programmer for software
Natasha Dillabough (MA): Database (over 3000 pieces)
Alexander Morgan (PhD): Programming, graphing

Music Technology

Ashley Burgoyne (PhD; now post-doc in Amsterdam)
Andrew Hankinson (PhD): Database
Gregory Burlet (MA): Database
Ryan Groves (MA): File conversion with Sikuli
Mike Winters (MA): Sonification, programming

Mathematics and Computer Science

Jamie Klassen (BA Mathematics and Statistics): Programming
Morgane Ciot (BA Computer Science and Linguistics): New Database
Saining Li (BA Computer Science and Linguistics): Programming

UK (Scotland: U of Aberdeen): Frauke Jürgensen, PI

Core team:

Prof. George Coghill (Computing Science, University of Aberdeen)
Dr Ian Knopke (BBC; Research Fellow, University of Aberdeen)

Collaborators:

Dr David Pearson (Psychology, University of Aberdeen)
Ralph Stelzenmüller (until 2013, PhD student at University of Aberdeen; director of Ensemble Combassal)

Advisors:

Prof. David Smith (Music, University of Aberdeen)
Dr Pang Wei (Computing Science, University of Aberdeen)

Student

Stephanie Colley (Research Assistant, University of Aberdeen)

US, New England (MIT and Yale): Michael Scott Cuthbert (MIT, Music), PI.

Ian Quinn (Yale, Music & Creative Consilience of Computing and the Arts), co-investigator.

Research Scientist Programmers:

Josiah Oberholtzer (PhD, Harvard University, Composition)
Ben Houge (MM, U Washington; now Asst. Prof. Berkelee College of Music)

Graduate Students:

Christopher White (PhD, Yale, Music Theory; now UNC Greensboro)
Kirill Zikanov (PhD, Yale, Musicology)

Undergraduate Programmers/*Co-Authors:

Maura Church* (AB, Harvard, Applied Math; now at Google)
Beth Hadley* (BS, MIT, Computer Science)
Erin Ibarra (BS, MIT, undeclared)
Lars Johnson* (BS, MIT, Computer Science)
Carl Lian (BS, MIT, Computer Science)
Evan Lynch (BS, MIT, Computer Science)
Daniel Manesh (BS/MS, MIT, Comp. Sci and Music)
Dylan Nagler (AB, Harvard, Comp. Sci. and Music)
Christopher Reyes* (AB, MIT, Comp. Sci/Elec. Eng. and Music)
Varun Ramaswamy (BS, MIT, Mathematics and Music)

EMMSAP (Electronic Medieval Music Score Archive Project):

Anna Grau Schmidt (PhD UPenn; Adj. Prof. DePaul University), project lead
Jeremy Jennings (MA, U. Oregon)
Kelsey Cowger (PhD, UCLA)
Corwyn Wyatt (MA, Boston U.)

Collaborating Researchers:[1]

Hans-Peter Kriegel (Prof. Ord., LMÜ Munich, Germany, Comp. Sci.)
Vladimir Viro (PhD, LMÜ Munich, Germany, Computer Science)
Matthias Röder (Director Herbert von Karajan Institute, Salzburg)
Christopher Ariza (PhD, NYU; formerly Visiting Asst. Prof., MIT)
Rob C. Wegman (Assoc. Prof., Princeton Univ., Music, data contributor)
Kris Shaffer (Asst. Prof., Charleston Southern University, Music Theory, formerly PhD., Yale)

Advisory Board for the whole team:

Cobb, Tom (L'Université du Québec à Montréal)
Freedman, Richard (Haverford College and CESR, Tours)
Gjerdingen, Robert (Northwestern University)
Headlam, David (Eastman School of Music and University of Rochester)
Huron, David (Ohio State University)
Pugin, Laurent (Répertoire International des Sources Musicales Bern, and University of Geneva)
Rodin, Jesse (Stanford University)
Smith, David (University of Aberdeen)

[1] Collaboration with the Germany and Austria-based researchers was funded by a separate grant, but the topic of research was identical to the ELVIS project.

3. Management, meetings and milestones for the project as a whole

Management

Julie Cumming, the PI of the McGill team, and the author of the original Digging into Data grant proposal, was the overall coordinator for the local teams. Most of our meetings were held at McGill (since McGill had the largest number of investigators, this was financially easier). CIRMMT (the Centre for Interdisciplinary Research in Music Media and Technology, housed at McGill) supported our meetings by hosting events in the form of CIRMMT workshops. These workshops, which involved non-team members and students, also served as a way of publicizing the project. McGill also hosted a website and a blog on the project. Ichiro Fujinaga, a co-investigator on ELVIS, had a SSHRC Partnership development grant while we had the ELVIS grant, and applied for a Partnership grant for the SIMSSA project (Single Interface for Music Score Searching and Analysis). The ultimate plan is for ELVIS technologies to become part of SIMSSA, as the "search and analysis" axis. As a result, Cumming and Fujinaga participated in a variety of events featuring both ELVIS and SIMSSA.

The McGill team had weekly meetings for most of the grant period; these were essential for working out our aims and approaches, and for assigning work to different members of the team, making decisions about priorities, developing research projects. Minutes from these meetings were posted on the original ELVIS website for the whole team.

The McGill team also maintained an ELVIS blog (http://simssa.ca/elvis, written by Catherine Motuz), that documented our progress and our achievements, while developments for the US team were posted to the music21 mailing list and blog (http://music21-mit.blogspot.com).

The Aberdeen held monthly meetings to monitor progress and determine short-term tasks and goals. Ian Knopke, being based in London rather than Aberdeen, connected to these meetings via Skype. In addition, Frauke Jürgensen met separately with Research Assistant Stephanie Colley, to go over technical details of tasks such as symbolic encoding and to devise a protocol for annotating our test sets of data for the module detecting tools. Ian Knopke came to Aberdeen at roughly six-monthly intervals for intensive work sessions lasting several days, during which he also presented talks, variously in the Music Research Seminar series and to groups of Computing Science students.

The US team was distributed between MIT and Yale but met together at meetings of the Northeastern Music Cognition Study Group, Society for Music Theory, at larger ELVIS meetings, and through e-mail. Code and data were shared via e-mail and through public Google Code and GitHub repositories. The US PI, Michael Cuthbert, wrote bi-annual reports to the NEH on their research progress, which the US team shared with the multinational project team.

Representatives from the UK and US teams travelled to Montreal twice during the course of the grant to attend multi-day meetings on which the various teams involved presented pieces of the overall puzzle, groundwork was laid for the mechanics of managing the investigative process, and a "hack day" was held to exchange ideas and help each other with specific problem-solving.

Team meetings and milestones

Each team meeting became a milestone, since they provided deadlines for development of tools and presentation of research.

Feb. 16, 2012, McGill. CIRMMT Workshop on processing large amounts of musical information. http://www.cirmmt.org/activities/workshops/research/musical-info/event

> This was the first meeting of the team as a whole, and allowed us to work out issues of how we would work together.

Nov. 2, 2012. New Orleans. Team meeting with members of the advisory committee. This meeting was held at the joint annual meeting of the American Musicological Society, the Society for Music Theory, and the Society for Ethnomusicology.

> Each team presented their work to the whole team, and we discussed methods and results. This was also the only meeting with members of the advisory committee, who provided insight and guidance into the project.

May 10-11, 2013. Workshop on searching symbolic musical data for repeated patterns: the ELVIS project.
http://www.cirmmt.org/activities/workshops/research/elvis/event
http://www.cirmmt.org/activities/workshops/research/elvis/elvis-2

> The first day was a workshop (May 10) the second day was a Hack Day (May 11). Team members and other users used the ELVIS online tools and discussed their research problems and solutions.

ELVIS blog post: http://simssa.ca/node/50

September 7, 2013. CIRMMT workshop: Workshop on SIMSSA (Single Interface for Music Score Searching and Analysis) III.

http://www.cirmmt.org/activities/workshops/research/simssa0913/event

ELVIS blog post: http://simssa.ca/node/80

October 12, 2013. Digging into Data Conference.

> Representatives from national teams, along with local student participants, met to discuss future plans. Collaboration at the conference resulted in the music21 tools for diad reduction for complex works.

Lessons and successes of the project

Different countries, different expectations. The differing practices and expectations **of** the different national granting agencies did lead to some problems. One was the actual available money. The US team had much less money to spend due to the amount of grant money that went to their institutions for indirect costs of research and subawards. The US and UK had greater reporting requirements during the grant period than Canada did. Timescales for various targets and the technicalities of administration were not always compatible (Canada finished December 31 2013; UK finished April 1, 2014; and US finished August 31, 2014). Once the differing structures of the academic year in the various countries were added in, communication between the various teams was not as consistent as we had hoped for. This was mitigated to a certain extent by the project-wide meetings in Montreal, and by the communication facilities offered through the main ELVIS website at McGill. However, for future projects, it would be good to include a more detailed communication plan. This might include more regularly-scheduled conference calls among the various PIs, so that contact is maintained even during the busiest times of the year.

The Aberdeen team found the administrative load especially weighty. The volume of the paperwork was daunting, as well as the style, which is unfamiliar to historical musicologists. We learned that in future grants, this administrative overhead needs to be considered more carefully, so that we can plan for the inclusion of an experienced part-time administrator so that the scholars can a greater proportion of their time on the research itself. Because MIT does not have a graduate program in music (and the CS students are well funded) finding a programmer able to work carefully with musical notation posed an initial problem for the US team necessitating a three-month extension in the project to cover a gap in programmer availability. The delay was, however, more than made up for by the quality of programming expertise eventually found.

All the teams found that they had to adjust their budgets as the actual expenses of running the grant became clearer, but fortunately this turned out to be possible in all cases, and we recommend that flexibility in this area be maintained.

Overall, however, the project has been extremely successful. We all benefited from working together and sharing our ideas, frustrations, and solutions. The visibility of the Digging into Data program increased the status of our work and of the Digital Humanities more generally within our institutions and our disciplines. Thanks to the grant, we have made huge advances in analyzing counterpoint with computers, provided resources and tools for other scholars, and trained many developing scholars who will also continue ouwork.

Part II: THE MCGILL TEAM

1. History of the project

Getting Started with the database (January to June 2012). Our first priority was to set up a database where we collect files of music in symbolic notation, with searchable metadata. We had already decided to use music21 as the foundation of our analysis tool. Fortunately music21 can parse a large variety of file types. In order to get up and running fast we used the Drupal Content Management system. Team meetings focused on establishing data entry guidelines and the structure of the database. Guidelines were adapted and modified as we uploaded more different genres and categories of pieces. The minutes of our meetings and the guidelines for data entry are on the database. PhD Student Catherine Motuz made a major contribution to the structure of the tags and the logic of the Database. PhD student Jacob Sagrans also contributed in this area.

Finding music in symbolic notation. As we developed the database we also searched for music already in symbolic notation to add to the database. We approached many different scholars, who contributed files to our project. Students working on the project uploaded files that they had created; we searched online sources such as Choralwiki and composer sites; we uploaded large portions of the data set used by the Yale team (largely 18th- and 19th-century music). We also approached a variety of scholars, who gave us access to their files. They include: Rob Wegman, fifteenth-century Masses; Jesse Rodin, the Josquin Research Project; Richard Freedman, files from the Du Chemin chansons; Clemens Goldberg, fifteenth-century chansons.

We continued adding files to the database until the end of the granting period. During the second year of the grant we became discontented with the limitations of the Drupal interface, and decided to change to a customized Django interface. This work is ongoing. MA student Natasha Dillabough uploaded more than half of the files to the database, making a major contribution to the project.

ELVIS blog posts.

- March 12th, 2012: Putting Drupal to good use. Describes the pros and cons of using an out-of-the-box platform like Drupal to organize a first ELVIS site and database. http://simssa.ca/node/51
- June 7th, 2012: Categories vs. Tags. How to organize data so that individual pieces are findable vs. how to organize data for corpus research. Traditional categories such as those used by the Library of Congress are contrasted with computer-based ways of understanding similarities, specifically tags. http://simssa.ca/node/66
- September 3rd, 2012: Building a Corpus. Mentions all the sources of our pieces and describes the wealth of musical sources available online. http://simssa.ca/node/69

- April 3rd, 2013: Batches of Data: CMME and Sikuli. Using outside-the-box methods for assembling and converting pieces for the ELVIS database. http://simssa.ca/node/76

VIS: contrapuntal Analysis tool (March to October 2012). Christopher Antila (an MA student in music theory) was our lead programmer; he worked closely with a terrific undergraduate mathematics major, Jamie Klassen. He first got up to speed on Python and music21, and then began to develop our tools. This involved extensive discussion with the faculty members on the McGill team, as we figured out what we could do and what we wanted to know in conjunction with the software development. By 26 June 2012 we had a basic program named VIS (Vertical Interval Successions) on GitHub that could search for vertical intervals and interval n-grams and sort them in terms of frequency of occurrence. By August we had a working desktop interface that could be installed on the computers of team members. The basic features of this interface is retained in the current desktop application, and in the Web application of 2013.

ELVIS blog posts:

- September 17th, 2012: Introducing VIS. Discusses the difficulty of bridging the gaps between musicologists and programmers and praising scholars with both music-theoretical knowledge and programming skills. http://simssa.ca/node/70
- October 12th, 2012: Queries! On figuring out how to formulate queries to a computer, and what kinds of queries can tell us something about musical style. http://simssa.ca/node/72

Meeting in New Orleans (November 2012). We did three presentations in New Orleans: one in a paper session on digital initiatives in early music; one at a lunch for musicologists and music theorists about Ichiro Fujinaga's SIMSSA project (Single Interface for Music Score Searching and Analysis); and a private presentation for the whole team and members of the advisory board. The presentations were a major milestone for us, where we did a live demo of the desktop application of VIS. We also began to use VIS to query files from our database, comparing the vertical intervals in music from different periods, and comparing n-grams in different masses and models. As we added more features to VIS, each component became too complicated for its own good.

ELVIS blog posts.

- November 7th, 2012: Everyone needs a Nate Silver. About how statistical models and the concept of entropy can be used to discuss musical style. http://simssa.ca/node/73
- November 2012: ELVIS at the AMS! http://simssa.ca/node/71

Rewrites of VIS (January to August 2013). We (Christopher Antila, Jamie Klassen, and Alexander Morgan) did two rewrites of VIS. We separated the interface from the

analysis components, and created a more modular structure that made it much easier to solve problems with the software, and that also made it easy to increase the functionality of the program.

VIS: modular structure and web application (September to December 2013). Problems with installing the VIS desktop application on different computers, plus the desire to make it easily available to users, with no need for them to do updates, etc., we decided to create a web application during the final months of the grant. We now have three programs: an analysis tool (the n-gram indexer), a desktop interface, and Web interface, including an API. Analysis tool: The n-gram indexer, which can coordinate and index two musical dimensions. Our focus has been on contrapuntal dimensions – vertical and horizontal intervals. However, it can also be used on other musical dimensions: pitch and duration, or bass and intervals in the upper voices (figured bass analysis).

2. Research

While our overall focus was in creating resources for the general public of music scholars we also used the database and our tools to carry out research throughout the project. For detailed discussion of our database and software, see the ELVIS website: http://elvisproject.ca/ and our GitHub site: https://github.com/ELVIS-project.

We have already made multiple presentations and there are forthcoming publications; our research work is ongoing. Research projects include:

Style change. The central issue for our project is style change, and all our research touched on this issue. A central project (see Antila and Cumming, 2015) took three different datasets of music from different periods (the mid-fifteenth century, the period around 1500, and the mid-to-late sixteenth century). Each dataset included music by different composers in different genres. We calculated the most common interval 3-grams for each set, and looked at the relative representations of three-grams across the three sets (we represented our findings in a Venn diagram of 3-gram clouds). We found that the most common 3-grams were common to all three datasets, and that each set shared some three grams with the earlier data set, but that there was a clear progression through time, with common 3-grams from the earliest dataset giving way to new 3-grams in the later set. We also found different degrees of repetition among the three sets. This was a group project, involving five or six graduate students and three faculty members.

Contrapuntal repetition. Repetition is an essential feature of almost all music, but the amount of repetition varies across time and genre. Exact quantification of repetition is almost impossible without computers. So we set out to quantify the repetition in a test set of the Duos by Orlandus Lassus from 1577 (see Schubert and Cumming, 2015). In order to do the calculations we had to decide what we meant by repetition, and discovered that there are three different kinds, and that looking at the most common

three-grams in different pieces provides insight into the characteristic patterns of individual works.

Theory and practice. Music theorists from all periods describe the music of their own time. However, it sometimes seem that what they describe does not fit the repertoire. Alex Morgan (see Morgan) decided to look at the treatise by Johannes Tinctoris on counterpoint (1477) in relation to music by composers whom Tinctoris praised in the treatise. Tinctoris includes a huge list of acceptable 2-grams in his thesis, and comments on their importance and frequency. Morgan compared the representation of those 2 grams in the repertoire to Tinctoris's commentary on the 2-grams, and found a high degree of correlation.

Sonification of music data. Michael Winters, a graduate student in music technology whose own research is on sonification of all kinds of data, joined the ELVIS project in the Fall of 2013. He worked on the VIS programming, and also developed a tool for sonifying data produced by VIS (see Winters and Cumming, 2015). The resultant sound files (for composers of different periods) sounded significantly different.

3. Student training

Student training was a central features of the McGill team's work. Seventeen different students worked on the project in a variety of roles. Only a few students worked for the whole duration of the grant (Christopher Antila, Andrew Hankinson, Catherine Motuz, and Jacob Sagrans, Natasha Dillabough), but there were always at least six students working on the project. The weekly meetings included the students on the project and the faculty co-investigators, and we all learned from each other, and figured out how to talk across the humanities science divide, as we developed research questions and projects, and figured out how to explain and structure those questions so that we could use the computer to answer them.

4. Presentations

Cumming, J., and I. Fujinaga. Member of a Panel discussion. "New Digital Projects for the Study and Dissemination of Medieval and Renaissance Music." American Musicological Society session organized by Richard Freedman and Mauro Calcagno. AMS/SEM/SMT annual meeting New Orleans, Nov. 1-4, 2012.

Morgan, A. "The Phrygian Cadence and Modal Identity in Victoria," College-Conservatory of Music, Cincinnati, USA, Feb. 23, 2013.

Cumming, J. "ELVIS." Montreal Digital Humanities Showcase. May 22, 2013.

Cumming, J., and P.Schubert. "Another Lesson from Lassus: Quantifying Contrapuntal Repetition in the Duos of 1577." Med-Ren Certaldo, July 7, 2013.

Morgan, A. "Testing Tinctoris: Computer-Assisted Analysis of the *Liber de arte contrapuncti*." Med-Ren Certaldo, July 7, 2013.

Cumming, J. "ELVIS: Electronic Locator of Vertical Interval Successions." Montreal, Digging into Data Challenge 2013 Conference, 12 October, 2013.

Antila, C. "Electronic Locator of Vertical Interval Successions." Montreal DH Showcase / Vitrine HN 2014. Université du Québec à Montréal, January 24, 2014. Received prize for best presentation.

Winters, M., and J. Cumming. "Sonification of Symbolic Music in the ELVIS Project." 20th International Conference on Auditory Display (ICAD–2014) June 22–25, 2014, New York, USA.

Antila, C., and J. Cumming. "The Vis Framework: Analyzing Counterpoint in Large Datasets." Poster at the *International Society for Music Information Retrieval*, Taipei, Taiwan, October 27-31, 2014.

Morgan, A. "Contrapuntal corpus studies on the *Liber de arte contrapuncti* and contemporaneous repertoire." Conference on Johannes Tinctoris and Music Theory in the Late Middle Ages and Early Renaissance, London, England. Oct. 9, 2014.

Morgan, A. "Testing Tinctoris: Data-Mining a Renaissance Treatise." Poster. Society for Music Theory Annual Meeting, Milwaukee, USA. Nov. 7, 2014.

Schubert, P., and J. Cumming. "Another Lesson from Lassus." Society for Music Theory Annual Meeting, Milwaukee, November 2014.

5. Publications (forthcoming)

Schubert, P., and J. Cumming, 2015. "Another Lesson from Lassus: Quantifying Contrapuntal Repetition in the Duos of 1577." *Early Music*.

Winters, R. M., and J. Cumming, 2014. "Sonification of Symbolic Music in the ELVIS Project." *Proceedings of the 20th International Conference on Auditory Display* (ICAD–2014) June 22–25, 2014, New York, USA.

Antila, C., and J. Cumming, 2015. "The Vis Framework: Analyzing Counterpoint in Large Datasets." *Proceedings of the International Society for Music Information Retrieval*, Taipei, Taiwan, October 27-31, 2014.

Part III. THE ABERDEEN TEAM

1. History of the project.

In April 2012, Frauke Jürgensen was invited to Study Days held in Graz, Austria, to consider research collaboration possibilities between the fields of study of Cognition and Early Polyphony. Participants were invited to submit collaborative papers to a special issue of the Journal of Interdisciplinary Music Studies. As a result of this, David Pearson (of the Department of Psychology at the University of Aberdeen) was invited to conduct a collaborative project in which computing techniques to assess the vertical interval content of collections of music were combined with a psychological experiment that attempted to measure the reactions of subjects exposed to the music in different ways. The resulting paper was accepted for publication, and will be appearing in the upcoming issue of the *Journal of Interdisciplinary Music Research*.

Throughout the course of the grant, the UK team presented results of their research at the Medieval and Renaissance Music Conference 2012 (Nottingham) and 2014 (Birmingham) (details in publications and presentations list), as well as at Research Seminars held in Aberdeen. A research paper emerging from the 2012 Medieval and Renaissance Music Conference presentation is in preparation. Towards the end of the project, we began discussions with colleagues in the biomedical sciences at the University of Aberdeen with the aim of designing a follow-up project. This is still in the conceptual stages.

Our Launch Event was in the context of an international conference entitled "Beyond the Semitone", held in October 2013 at the University of Aberdeen and organised by Frauke Jürgensen in collaboration with composer Geoff Palmer. This conference attracted approximately 40 participants and invited performers and speakers, from across the globe. The conference crossed musical disciplines that have an interest in studying questions of tuning, chromaticism, and microtonality throughout music history: there were historial musicologists and ethnomusicologists, composers, performers, musical instrument designers, and physicists, participating in concerts, academic papers, and composition workshops. ELVIS was featured in two central events of the conference. First, Ensemble Combassal (dir. Ralph Stelzenmüller) gave a semi-liturgical-reconstruction performance of Carlo Gesualdo's Tenebrae Responses for Good Friday, which was open to the public at large as well as conference attendees and university members. Carlo Gesualdo was chosen for a happy co-incidence of reasons: 2013 was the 400th anniversary of his death, and he is especially known for his daring use of chromaticism, leading to especially unusual successions of vertical intervals. This set the stage for an extended lecture-demonstration at which questions surrounding the stylistic context and performance practice of the work were raised, and where the contributions of the ELVIS project for studying such questions were elucidated. This event was again open to the general public, as well as the conference audience. A special issue of the electronic journal Scottish Music Review, dedicated to this

conference, is being edited, and a research paper detailing our contribution is in submission to this issue.

2. Research

The UK team in particular was interested in studying cases of repeated vertical interval successions where there were inexact matches. This is the case with so-called modules, compositional building blocks that can be varied through contrapuntal transformations, ornamentation, or other operations. We devised an annotation system based on Peter Schubert's analyses-by-hand of Palestrina's first book of four-voice motets, which allows us to test our module-detection algorithms against a ground truth. Beyond methodological developments, we applied pre-existing, simpler techniques to the existing datasets to make comparisons among Renaissance and early Baroque composers, showing that stylistic differences can successfully be described using vertical interval successions.

Our approach to identifying contrapuntal modules originally was based on the use of suffix trees. While this worked for finding exact melodic and harmonic matches for entire collections, such as the complete set of Palestrina Masses, it was not effective in locating ornamented melodic material. From an algorithmic point of view this can be viewed as specialised form of an approximate string matching problem. After exploring various edit-distance and n-gram methods, we have developed an approach to locating similar contrapuntal modules based on a combination of three techniques. First, we expanded our original suffix tree based methods to use suffix arrays, which are a more compact version of suffix tress that have some additional desirable search properties. With windowing, suffix arrays are basically a very efficient form of n-gram search. Secondly, there have been some relatively recent advances in suffix array algorithms towards approximate search techniques that have proven to be effective in our work. Finally, we have also combined these with a new music encoding scheme we have called Interleaved Melodic Encoding, that can be used to encode multidimensional (pitch, duration) attributes into traditional string pattern-matching algorithms, while still maintaining the approximate matching techniques within the suffix-array algorithm. The combination of these three techniques has proven extremely effective for our particular problem.

To evaluate the effectiveness of our approach we have adopted a three-part strategy. First, we hand-encoded a number of specialised cases as an initial evaluation set, and to work out basic elements of the technique. We also encoded a number of pieces from the Palestrina motets, where modules had been located by musicologists and made sure we could reproduce these results. Working through these two initial evaluation sets then gave us the confidence to apply these techniques to the larger complete set of Palestrina Masses, where the materials would be far too large to ever check every instance by hand. We are currently working these results into a submittable journal paper, as well

as exploring adequate methods of visualising large sets of contrapuntal modules and the underlying structure they expose.

3. Techniques

Our primary toolkit is the PerlHumdrum toolkit for Symbolic Music Processing. This is basically a collection of tools for reading, writing, and processing musical scores that have been computationally-encoded in a symbolic form, originally aimed at Humdrum scores. The toolkit has continued to evolve based on our research needs, and this project has led to many enhancements, mostly in the area of contrapuntal analysis. Another change is the use of some new tools in the Haskell computing language, and more enhancements to the accompanying PerlLilypond package for displaying scores.

Other work has mostly been in the area of statistical tools that operate across entire collections. The temperament detection work has necessitated the building of models of various traditional temperaments, such as Just, Pythagorean, and Mean-Tone models. Once constructed, these are treated as probability models. We then build harmonic models of both individual and collections of pieces and compared against temperament distributions for best matches. Similar techniques were used for the "Authentic Listener" work, but based on consonance and dissonance models.

4. Presentations

Jürgensen, F. "Teaching keyboard improvisation in the Renaissance: Memory and the Fundamenta of Paumann and Buchner," ESF Exploratory Workshop on "Cognition of Early Polyphony,"Graz, Austria (April 2012).

Jürgensen, F. and I. Knopke. "Extracting Possible Temperaments from Historical Scores using Symbolic Means." Medieval and Renaissance Music Conference, Nottingham, UK (July 2012).

Jürgensen, F. with R. Stelzenmüller and Ensemble Combassal. Lecture recital, "*Iternarium extaticum celeste* (Kircher, Rome, 1656): An ecstatic journey through the experimental world of polytonality in 17th-century Italy." Beyond the Semitone, Aberdeen, UK (October 2013).

Jürgensen, F. "Interdisciplinary approaches to the history of musical style in the Renaissance: the ELVIS project." Centre for Early Modern Studies, Aberdeen, UK (March 2014).

Jürgensen, F., and I. Knopke. "Computerised style classification of late Renaissance polyphony." Medieval and Renaissance Music Conference, Birmingham, UK (July 2014).

5. Publications

Jürgensen, F., D. Pearson, and I. Knopke, 2014. "Building an Authentic Listener: applying a passive exposure-based training paradigm to detecting difference among compositional styles," *Journal of Interdisciplinary Music Research* (26 pages) In press.

Jürgensen, F., and I. Knopke, 2015. "Counterpoint and Chromaticism: the intersection of compositional style and the theory of chromatic music in the *Tenebrae* of Carlo Gesualdo," *Scottish Music Review*. In submission.

Jürgensen, F., and I. Knopke. "Empirical Determinations of Tuning Systems in Fifteenth-Century Keyboard Music." In preparation.

Part IV. THE US TEAM

1. History of the Project

The US team's Investigators (Cuthbert and Quinn) had been working together on ways to improve computational musicology and music theory since early 2011. Since 2006, Cuthbert had been developing music21, a toolkit for computational musicology. Based on the principles of flexibility and musical vocabulary developed by David Huron's "Humdrum" toolkit from the 1980s and early 1990s, music21 added a modern, object-oriented framework and the ability to work with scores encoded in any of nearly a dozen common music notation formats. By the start of the project, music21 was able to aid researchers in solving nearly all of the common problems that needed to categorize and analyze vertical interval successions, but the toolkit was not optimized to work with corpora of more than a few hundred compositions.

Quinn was already recognized as a pioneer in mathematical music theory, but also felt that statistical and data-based approaches to musical analysis had great future potential. At Yale, he had attracted a cohort of students with both the mathematical and musical fluency to tackle tough problems in corpus analysis.

2. Research and Techniques

Quinn began by confronting a fundamental problem posed by our collective project, "What would an electronic locator of vertical interval successions look like?" He developed a concordancer program, built on music21, for examining the frequency of four-note pairs and larger groupings of vertical successions in the chorales of J.S. Bach. His research showed that certain commonly taught rules about tonal harmony were actually exceptional cases in the music of Bach.

Quinn's student, Christopher Wm. White, and Cuthbert approached the problem of reducing the size of corpora to be analysed from two different angles. White used an alphabet reduction method in music21 applied to the tonal MIDI corpus. The number of distinct vertical interval successions in, say, a nineteenth-century opera, or 200 symphonies composed between 1750 and 1800, is too vast to get meaningful signal from the noise. White's software classified each vertical sonority and then, using an algorithm modified from spell checkers, folded the least common sonorities into more common ones. Through his iterative process, he eventually reduced the vocabulary for each fifty-year period to approximately one dozen chords and one hundred vertical successions. His work showed that there were substantial changes from one period to the next, and that the chords and successions generally thought to be fundamental to musical languages were not in fact the ones used most commonly. White and Quinn have posted data on individual composers on the Yale-Classical Archives Corpus on line (ycac.yale.edu); see White and Quinn presentation, 2014.

Cuthbert, working with undergraduate researchers at MIT and Josiah Oberholtzer, developed systems to identify less salient musical moments using conventional music theoretical knowledge of passing tones, neighbour tones, etc., and thereby reduce a piece to its voice-leading skeleton.

Work on music21 focused on eliminating inefficiencies ($O(n^2)$ or $O(n!)$) in the analysis of larger scores, largely through an AVL-Tree method for efficiently retrieving musical elements in large scores and score collections. Music21 was also redesigned to work effectively in distributed computing environment (Cuthbert and Hadley) and, through a published API, enabled computational musicology techniques, such as motivic searches, within other, non-score-based music projects (Cuthbert and Johnson). Work on making music21 easier for beginning programmers and musicians with no programming experience also progressed, resulting in real-time score interaction with large datasets (Cuthbert and Reyes), which would become refined in IPython Notebook versions and a JavaScript port, music21j.

Cuthbert also pursued the use of music21 in analyzing similarities (vertical and melodic) between musical works in the fourteenth and early fifteenth-centuries. "Monks, Manuscripts, and Other Peer-to-Peer Song Sharing Networks of the Middle Ages" uses statistical methods borrowed from the analysis of peer-to-peer file sharing to show how most existing models for musical transmission in the Middle Ages are highly unlikely from a statistical standpoint. Chapter 4 of his monograph identifies two previously unknown works of fourteenth-century music using motivic similarity techniques and a large dataset of medieval music developed as part of the ELVIS project, the EMMSAP repository.

Through collaborations established with a separate grant for computational musicology exchanges with Germany, the US team gained access to the data and tools for working with the IMSLP dataset of 500GB of symbolic data for over 100,000 full scores. This data was created through optical music recognition (OMR) and unfortunately our tests showed that it was too low quality for statistical analysis of any but the crudest tests. Maura Church and Cuthbert developed a system for reducing the number of rhythmic errors in the system by 20% in each part: for a six-part score, this results in a 75% reduction of total errors. The US team is continuing their ELVIS/DiD funded research until the end of the summer and plans to apply these techniques to the complete IMSLP dataset.

3. Presentations

Cuthbert, Michael Scott:

"Searching and Analyzing Renaissance Musical Scores with Computational Methods," Renaissance Society of America, New York City, 27 March 2014

"Music, Manuscripts, and Megabytes: Unlocking sound in the *Ars Mutandi* (1340–1420)," MIT Ancient and Medieval Studies Colloquium Series, 19 September

"Identifying Quotations and Concordances in Ars Nova Music with Computational Methods," *Medieval and Renaissance Music Conference,* Certaldo, Italy, 4–9 July 2013.

"The Credo 'Qui sonitu melodie': Quotation, Paraphrase, and Influence in the *Ars Mutandi*," invited lecture, *International Symposium on Late Medieval and Early Renaissance Music,* Kloster Neustift/Novacella, Brixen, Italy, 26 June 2013.

"`music21`, EMMSAP, and ELVIS: New Connections in Searching and Reduction of Large Corpora," annual meeting, Center for Interdisciplinary Research in Music Media and Technology, McGill University, 10 May 2013.

"The Right Tune in a Sea of Notes: Searching Musical Scores with Musical Intelligence," invited lecture, Harvard University forum, *Expanding the Boundaries of Authorship,* 15 April 2013.

"With Computer, Ear, and Mind: New Approaches to Centuries-Old Problems in Music History," invited lecture, Dartmouth College, 7 March 2013.

"Hacking Symbolic Music with the `music21` Toolkit," presentation, inaugural meeting, European Forum for Digital Musicology co-sponsored by the Vienna Classical Music Hack Day, Vienna, 2 February 2013.

"Humdrum and `music21`: New and successful approaches to Digital Musicology," invited lecture, Ohio State University, 21 November 2012.

"Teaching Musicology in a 2.0 World," invited presentation for the "Master Teacher Session: Teaching the Discipline, Disciplining the Teacher" at AMS, New Orleans, 1–4 November 2012.

"Common Blunders and Golden Opportunities: Learning from Past Digital Music Projects," Conference Paper, *Transcribing the Beneventan Chant*, Radcliffe Institute for Advanced Study, 19 October 2012.

"Understanding Music in the Age of Plague and Schism with Computer, Mind, and Ear," invited talk, Distinguished Lecture Series, UCLA Music Department. 11 October 2012.

"Digital Musicology of Late Medieval Polyphony," Radcliffe Institute for Advanced Study, 19 September 2012.

"Codieren von Musiknoten für analytische Abfragen," invited lecture, Department of Computer Science, Ludwig Maximilians Üniversität, Munich, 12 July 2012.

"Corpus Research using the `music21` Toolkit," keynote lecture, Northeast Music Cognition Group semiannual meeting, Yale University, 28 April 2012.

"What were the odds?: Reexamining Early (and not-so-Early) Music with statistical models," invited lecture by the graduate students, Department of Music, Duke University, 30 March 2012.

"Understanding Musical Corpora with `music21`: Completed Tasks and Future Applications," Workshop on Processing Large Amounts of Musical Information, McGill University, 17 February 2012.

White, Christopher Wm.:

"Style, Similarity, and Tonal Communication in the Common Practice," Washington University in St. Louis, Dept. of Music Lecture Series. 2013.
"Metric Induction using Autocorrelation, Templates, and Probabilities," Workshop on Searching Symbolic Musical Data for Repeated Patterns, McGill University, 2013
"Computational Music Research at Yale University," Northeast Music Informatics Special Interest Group, Dartmouth College, 2012.
"HMMs and Musical Syntax," Computational Linguistics Reading Group, Yale University, 2012.

White, Christopher Wm., and Ian Quinn. "Compiling and Processing the Yale-Classical Archives Corpus." International Conference on Music Perception and Cognition, August 2014 (Seoul).

Quinn, Ian:

"A Corpus-Based Approach to Tonal Theory," Society for Music Theory, Workshop Session, 2012.

4. Publications

Cuthbert, Michael Scott, "Monks, Manuscripts, and Other Peer-to-Peer Song Sharing Networks of the Middle Ages," in *Cantus scriptus: Technologies of Medieval Song, 3rd Lawrence J. Schoenberg Symposium on Manuscript Studies in the Digital Age,* edited by Lynn Ransom (Piscataway, N.J.: Gorgias Press, 2012), pp. 101–23.

Cuthbert, Michael Scott, Beth Hadley, Lars Johnson, and Christopher Reyes "Interoperable Digital Musicology Research via `music21` Web Applications," *Proceedings of the Joint CLARIN-D/DARIAH Workshop, Service-oriented Architectures (SOAs) for the Humanities: Solutions and Impacts at the Digital Humanities Conference,* Hamburg, Germany, July 2012.

White, Christopher Wm., "An Alphabet Reduction Algorithm for Chordal *N*-grams," in *Mathematics and Computation in Music*, ed. J. Yust and J. Wild (Heidelberg: Springer, 2013).

White, Christopher Wm., "Some Statistical Properties of Tonality, 1650–1900," Ph.D. Dissertation, Yale University, December 2013. Ian Quinn, advisor.

White, Christopher Wm., "Changing Styles, Changing Corpora, Changing Tonal Models," in *Music Perception* (special issue on corpus analysis) 31.3 (2014): 244-53.

Church, Maura and Michael Scott Cuthbert, "Improving Rhythmic Transcriptions via Probability Models Applied Post-OMR," *Proceedings of the International Society for Music Information Retrieval* 15 (2014).

Cuthbert, Michael Scott, et al., `music21` versions 1.1 to 1.9 (2.0 to be released before the end of the US team's grant period). Reviewed: Dmitri Tymoczko, *Music Theory Online* 19.3 (September 2013); Ian Quinn, "Digital and Multimedia Scholarship," *Journal of the American Musicological Society* 67.1 (2014): 295-307.

In Press, Under Consideration, and Beta Software

Cuthbert, Michael Scott. `music21j`: A Web-based Toolkit for Musical Analysis and Visualization, v. 0.2.

Cuthbert, Michael Scott. *Ars Mutandi: Italian Sacred Music in the Age of Plague and Schism*, monograph, 89,000 words.

Quinn, Ian, and Christopher Wm. White. 2014. "Corpus-Derived Key Profiles Are Not Transpositionally Equivalent." Submitted to *Music Perception*.

White, Christopher Wm., and Ian Quinn. 2014. "Content-Free Models of Harmonic Function." In preparation.